

PERIOD ESTIMATION FOR SPARSELY-SAMPLED QUASI-PERIODIC LIGHT CURVES APPLIED TO MIRAS

SHIYUAN HE¹, WENLONG YUAN², JIANHUA Z. HUANG¹, JAMES LONG¹ & LUCAS M. MACRI^{2,3}

¹Department of Statistics, Texas A&M University, College Station, TX, USA

²George P. and Cynthia W. Mitchell Institute for Fundamental Physics & Astronomy, Department of Physics & Astronomy,
 Texas A&M University, College Station, TX, USA

³Corresponding author, lmacri@tamu.edu

ABSTRACT

We develop a non-linear semi-parametric Gaussian process model to estimate periods of Miras with sparsely-sampled light curves. The model uses a sinusoidal basis for the periodic variation and a Gaussian process for the stochastic changes. We use maximum likelihood to estimate the period and the parameters of the Gaussian process, while integrating out the effects of other nuisance parameters in the model with respect to a suitable prior distribution obtained from earlier studies. Since the likelihood is highly multimodal for period, we implement a hybrid method that applies the quasi-Newton algorithm for Gaussian process parameters and search the period/frequency parameter over a dense grid.

A large-scale, high-fidelity simulation is conducted to mimic the sampling quality of Mira light curves obtained by the M33 Synoptic Stellar Survey. The simulated data set is publicly available and can serve as a testbed for future evaluation of different period estimation methods. The semi-parametric model outperforms an existing algorithm on this simulated test data set as measured by period recovery rate and quality of the resulting Period-Luminosity relations.

Keywords: methods: statistical – stars: variables: Miras

1. INTRODUCTION

The determination of reliable periods for variable stars has been an area of interest in astronomy for at least four centuries, since the discovery of the variability of Mira (*o* Ceti) by Fabricius in 1596 and the first attempts to determine its period by Holwarda & Bouillaud in the mid-1600s. The availability of electronic computers for astronomical research half a century ago enabled the development of many algorithms to estimate periods quickly and reliably, such as [Lafleur & Kinman \(1965\)](#); [Lomb \(1976\)](#); [Scargle \(1982\)](#).

The aforementioned algorithms work best in the case of periodic variations with constant amplitude and Mira variables present several challenges in this regard. While their periods of pulsation are stable except for a few intriguing cases ([Templeton et al. 2005](#)), Mira light curves can exhibit widely varying amplitudes from cycle to cycle (see, for example, the historical light curve of Mira compiled by [Templeton & Karovska 2009](#)). In the case of C-rich Miras, the stochastic changes in mean magnitude across cycles (e.g., [Marsakova 1999](#)) only complicate the problem further. The wide variety of light curves for long-period variables, already recognized by [Campbell \(1925\)](#) and [Ludendorff \(1928\)](#), may com-

plicate the identification of Miras among other stars. Lastly, from a purely practical matter, it is simpler to obtain light curves spanning several cycles for RR Lyraes or Cepheids (with periods ranging from ~ 0.5 to ~ 100 d) than for Miras (with periods ranging from ~ 100 to ~ 1500 d).

Despite these challenges, the identification and determination of robust periods for Miras — especially in the regime of sparsely-sampled, low signal-to-noise light curves — would very beneficial for the determination of distances to galaxies of any type. Thanks to the unprecedented temporal coverage of the Large Magellanic Cloud (LMC) by microlensing surveys, the availability of large samples of extremely well-observed Miras has led to a thorough characterization of their period-luminosity relations at various wavelengths ([Wood et al. 1999](#); [Ita et al. 2004](#); [Soszynski et al. 2007](#)). The dispersion of the *K*-band period-luminosity relation ([Glass & Evans 2003](#), $\sigma = 0.13$ mag), is quite comparable to that of Cepheids at the same wavelength ([Macri et al. 2015](#), $\sigma = 0.09$ mag) and makes them competitive distance indicators.

The third phase of the OGLE survey ([Udalski et al. 2008](#)) imaged most of the LMC with little interruption

over 7.5 years and resulted in the discovery of 1663 Miras (Soszyński et al. 2009) with a median of 466 photometric measurements per object. The temporal sampling of these light curves and their photometric precision is exceptional relative to typical astronomical surveys and makes period estimation relatively easy. In comparison, a similar span of observations of M33 by the DIRECT (Macri et al. 2001) and M3SSS projects (Pellerin & Macri 2011) in the I band consists of a median number of 44 somewhat noisy measurements, heavily concentrated in a few observing seasons. Representative Mira light curves from the OGLE & DIRECT/M3SSS surveys are shown in Fig. 1. There are several reasons for the striking difference in quality between these two datasets. The LMC Miras are among the brightest objects in the OGLE fields, whereas their M33 counterparts are among the faintest in the aforementioned surveys of this galaxy. While the effective exposure times of all these surveys are quite comparable, after taking into account differences in collecting area of their respective telescopes, M33 lies approximately 6.2 mag farther in terms of its I -band apparent distance modulus. Furthermore, the main goal of the OGLE project (detection of microlensing events) requires a very dense temporal sampling of the survey fields; this is achieved by using a dedicated telescope and is helped by the fact that the LMC is observable nearly all year long from the site. In contrast, the observations of M33 were carried out using shared facilities (available only a few nights per month) with the primary purpose of studying Cepheids and eclipsing binaries (which do not require exceptionally dense temporal sampling), and the galaxy is only observable all night long for $\sim 1/3$ of the year. Standard period estimation algorithms which work well for high signal-to-noise, well sampled light curves such as those obtained by OGLE will fail on more typical datasets represented by the M33 observations. The purpose of this work is to develop and test methodology for estimating periods for sparsely-sampled, noisy, quasi-periodic light curves such as those of Miras observed in M33 by the aforementioned projects.

The rest of the paper is organized as follows: In §2 we review several existing period estimation methods. In §3 we introduce a new semi-parametric model for Mira variables which uses a Gaussian process to account for deviations from strict periodicity. We use maximum likelihood to estimate the period and the parameters of the Gaussian process, while other nuisance parameters in the model are integrated out with respect to some prior distributions using earlier studies. Since the likelihood is highly multimodal for the period/frequency parameter, we implement a hybrid method that applies the quasi-Newton algorithm for Gaussian process parameters and




Figure 1. Representative Mira light curves observed by OGLE-III in the Large Magellanic Cloud (top) and DIRECT/M3SSS in M33 (bottom).

a grid search for the period/frequency parameter. In order to assess the effectiveness of the semi-parametric model, in §4 we carefully construct a simulated data set by fitting smooth functions to the light curves of well-observed OGLE LMC Miras and resampling them at the cadence, noise level and completeness limits of the aforementioned M33 observations. Using the simulated data, in §5 we compare performance of existing period estimation methods to our semi-parametric model. We find that our proposed model shows an improvement over the generalized Lomb-Scargle model under various metrics. In §6, we conclude and discuss some future application. Simulated light curves for reproducing the results in the paper and performance benchmarking are made publicly available as supplementary material.

2. PERIOD ESTIMATION TECHNIQUES

Let y_i be the magnitude of a variable star observed at time t_i (in units of days) with uncertainty σ_i . The data set for this object, obtained as part of a time-series survey with n epochs is $\{(t_i, y_i, \sigma_i)\}_{i=1}^n$. One common approach to estimate the primary frequency of such an object is to assume some parametric model for brightness variation and then use maximum likelihood to estimate parameters. Zechmeister & Kürster (2009) define

the Generalized Lomb-Scargle (GLS) model as

$$y_i = m + a \sin(2\pi f t_i + \phi) + \sigma_i \epsilon_i, \quad (1)$$

where $\epsilon_i \sim \mathcal{N}(0, 1)$, m is the mean magnitude, a is the amplitude, $\phi \in [-\pi, \pi]$ is the phase, and f is the frequency (see [Reimann 1994](#), for early work in this model). Using the sine angle addition formula and letting $\beta_1 = a \cos(\phi)$ and $\beta_2 = a \sin(\phi)$ one obtains

$$y_i = m + \beta_1 \sin(2\pi f t_i) + \beta_2 \cos(2\pi f t_i) + \sigma_i \epsilon_i. \quad (2)$$

The likelihood function of this model is highly multimodal in f . However at a fixed f the model is linear in the parameters (m, β_1, β_2) . These two facts motivate the computation strategy of performing a grid search across frequency and minimizing a weighted least squares

$$\begin{aligned} &(\hat{m}(f), \hat{\beta}_1(f), \hat{\beta}_2(f)) \\ &= \arg \min_{m, \beta_1, \beta_2} \sum_{i=1}^n \frac{1}{\sigma_i^2} \{y_i - m \\ &\quad - \beta_1 \sin(2\pi f t_i) - \beta_2 \cos(2\pi f t_i)\}^2, \end{aligned} \quad (3)$$

at every frequency f on the grid. Under the normality assumption, the weighted least squares minimization is equivalent to maximizing the likelihood. Since the model is linear, computation of $\hat{m}(f), \hat{\beta}_1(f), \hat{\beta}_2(f)$ is straightforward. The residual sums of squares at f is

$$\begin{aligned} \text{RSS}(f) &= \sum_{i=1}^n \frac{1}{\sigma_i^2} \{y_i - \hat{m}(f) \\ &\quad - \hat{\beta}_1(f) \sin(2\pi f t_i) - \hat{\beta}_2(f) \cos(2\pi f t_i)\}^2, \end{aligned} \quad (4)$$

and the maximum likelihood estimator for f is

$$\hat{f} = \arg \min_f \text{RSS}(f). \quad (5)$$

Define RSS_0 as the (weighted) sum of squared residuals when fitting a model with only an intercept term m . The periodogram is defined as

$$S_{LS}(f) = \frac{(n-3)(\text{RSS}_0 - \text{RSS}(f))}{2\text{RSS}(f)}. \quad (6)$$

The periodogram has the property that if the light curve of the star is white noise (i.e., $y_i = m + \epsilon_i$), $S_{LS}(f)$ has an $F_{2, n-3}$ distribution. Thus the periodogram may be used for controlling the “false alarm probability,” the potential that a peak in the periodogram is due to noise ([Schwarzenberg-Czerny 1996](#)).

A large number of period estimation algorithms in astronomy are closely related to GLS. The Lomb-Scargle method is identical to GLS but first normalizes magnitudes to mean 0 and does not fit the m term ([Lomb 1976](#); [Scargle 1982](#)). The “harmonic analysis of variance” includes an arbitrary number of harmonics in Equation (1) ([Schwarzenberg-Czerny 1996](#); [Quinn & Thomson 1991](#)). [Bretthorst \(1988\)](#) incorporates Bayesian priors on the




Figure 2. Light curve of a Mira in the LMC observed by OGLE (black points), decomposed following Eqn. 7. Top panel: fitted light curve; middle panel: periodic signal, $m + q(t)$; bottom panel: stochastic variations, $m + h(t)$.

parameters β_1 and β_2 . The method is similar to performing a discrete Fourier transform and selecting the frequency which maximizes the [Deeming \(1975\)](#) periodogram. However, [Reimann \(1994\)](#) showed that GLS has better consistency properties than the Deeming periodogram.

It is also possible to use non-sinusoidal models but compute and minimize the residual sum of squares as above. For example, [Hall et al. \(2000\)](#) considers the Nadaraya-Watson estimator and [Reimann \(1994\)](#) uses the Supersmoother algorithm. [Wang et al. \(2012\)](#) uses Gaussian processes with a periodic kernel and found the period with maximum likelihood or minimum leave-one-out cross-validation error.

None of the above methods account for the non-periodic variation present in Miras. While these methods are adequate for densely-sampled Mira light curves (where the quantity of data overwhelms model inadequacy), their performance deteriorates in the sparsely-sampled regime. In Section 5, we compare our proposed model with the Lomb-Scargle method.

3. THE SEMI-PARAMETRIC MODEL

Suppose the data $\{(t_i, y_i, \sigma_i)\}_{i=1}^n$ is modeled by

$$y_i = g(t_i) + \sigma_i \epsilon_i,$$

where $g(t_i)$ is the light curve signal and the $\epsilon_i \sim N(0, 1)$ is independent with other ϵ_j 's. The signal of the light curve is further decomposed into three parts,

$$\begin{aligned} g(t) &= m + q(t) + h(t) \\ &= m + \beta_1 \cos(2\pi ft) + \beta_2 \sin(2\pi ft) + h(t), \end{aligned} \quad (7)$$

where m is the long-run average magnitude, $q(t) = \beta_1 \cos(2\pi ft) + \beta_2 \sin(2\pi ft)$ with frequency f is the exactly periodic signal, and $h(t)$ is the stochastic deviation from a constant mean magnitude, caused by the formation and destruction of dust in the cool atmospheres of Miras. Fig. 2 provides an example of the decomposition for a Mira light curve. The first two terms $m + q(t)$ in Eqn. 7 are exactly the same as the generalized Lomb-Scargle model of Eqn. 2. To simplify notation, we define $\mathbf{b}_f(t) = (\cos(2\pi ft), \sin(2\pi ft))^T$, so that $q(t) = \mathbf{b}_f(t)^T \boldsymbol{\beta}$. The subscript in $\mathbf{b}_f(t)$ emphasizes that the basis is parameterized by the frequency f .

A semi-parametric (SP) statistical model is constructed in Eqn. 7 if we assume $h(t)$ is a smooth function that belongs to a reproducing kernel Hilbert space \mathcal{H} with norm $\|\cdot\|_{\mathcal{H}}$ and a reproducing kernel $K(\cdot, \cdot)$. For this model, if the frequency f is known, we obtain a least squares kernel machine considered in Liu et al. (2007). Because the frequency is unknown, the response function is non-linear in f . This non-linearity and the multimodality in f of the residual sum of squares provides additional challenges that require a novel solution.

Besides the additive formulation in Eqn. 7, another possible solution to account for the quasi-periodicity is a multiplicative model such as $g(t) = m + h(t)q(t)$, where the amplitude of the strictly periodic term $q(t)$ is modified by a smooth function $h(t)$. However, the multiplicative model is more computationally intensive in nature and requires imposing a positive constraint on $h(t)$. As we will show in the following subsections, the $h(t)$ term in the additive model can be easily absorbed into the likelihood function. Nevertheless, the multiplicative approach is an interesting alternative approach to model formulation and is open to future study.

3.1. Equivalent formulations

Following §5.2 of Rasmussen & Williams (2005), for fixed f , the parameters m, β_1, β_2 and $h(t)$ in Eqn. 7 are jointly estimated by minimizing

$$\begin{aligned} \sum_{i=1}^n \frac{1}{\sigma_i^2} [y_i - m - \beta_1 \cos(2\pi ft_i) \\ - \beta_2 \sin(2\pi ft_i) - h(t_i)]^2 + \lambda \|h(\cdot)\|_{\mathcal{H}}^2, \end{aligned} \quad (8)$$

where λ is a regularization parameter. A smoothing/penalized spline model for $h(t)$ is a special case of the general formulation of Eqn. 8 with a specifically defined kernel; see §6.3 of Rasmussen & Williams (2005). For fixed λ , the solution of $h(t)$ is a linear combination of n basis functions $K(t_i, t)$, $i = 1, 2, \dots, n$, by the representer theorem (Kimeldorf & Wahba 1971; O'Sullivan et al. 1986). It is still left for us to choose the regularization parameter λ to balance data fitting and the smoothness of the function $h(t)$.

An equivalent point of view to the above regularization approach is to impose a Gaussian process prior on the function $h(t)$; see §5.2.3 of Rasmussen & Williams (2005). The benefit of this view is that it provides an automatic method for selecting the regularization parameter λ . In particular, we can absorb λ into the definition of the norm $\|\cdot\|_{\mathcal{H}}$ and assume the term $h(t)$ in Eqn. 7 follows a Gaussian process, $h(t) \sim \mathcal{GP}(0, k_{\boldsymbol{\theta}}(t, t'))$, with the squared exponential kernel $k_{\boldsymbol{\theta}}(t, t') = \theta_1^2 \exp\left(-\frac{(t-t')^2}{2\theta_2^2}\right)$, and parameters $\boldsymbol{\theta} = (\theta_1, \theta_2)$. The Gaussian process assumption implies that at any finite number of time points t_1, t_2, \dots, t_s , the vector $(h(t_1), \dots, h(t_s))$ is multivariate normally distributed, with zero mean and covariance matrix $\mathbf{K} = (k(t_i, t_j))$. This imposes a prior on the function space of $h(t)$. We also impose priors on m and $\boldsymbol{\beta}$ in Eqn. 7. In particular, we assume $m \sim \mathcal{N}(m_0, \sigma_m^2)$ and $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, \sigma_b^2 \mathbf{I})$. The prior mean m_0 can be interpreted as the average magnitude of Miras in a certain galaxy, and σ_m^2 is the variance of Miras in that galaxy; the prior variance σ_b^2 is the variance of the light curve amplitude. These prior parameters can be determined using previous studies. For example, in §5, we use well-sampled light curves of LMC Miras (Soszyński et al. 2009) to obtain values of these parameters. It is advisable to check the sensitivity of these prior specifications.

The benefit of using priors on m and $\boldsymbol{\beta}$ is three-fold: first, they introduce regularization by using information from early studies; second, they provide a natural device for separating the estimation of frequency and the light curve signal component using Bayesian integration when the parameter of interest is the frequency; lastly, the regularization parameter $\boldsymbol{\theta}$ of the non-parametric function is allowed to be chosen by the maximum likelihood, without resorting to the computationally expensive cross-validation method.

In summary, we have built the following hierarchical model for a Mira light curve:

$$\begin{aligned} y_i | m, \boldsymbol{\beta}, g(t_i) &\sim \mathcal{N}(g(t_i), \sigma_i^2), \\ g(t) &= m + \mathbf{b}_f(t)^T \boldsymbol{\beta} + h(t), \\ m &\sim \mathcal{N}(m_0, \sigma_m^2), \quad \boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, \sigma_b^2 \mathbf{I}), \\ h(t) | \boldsymbol{\theta} &\sim \mathcal{GP}(0, k_{\boldsymbol{\theta}}(t, t')), \end{aligned} \quad (9)$$

where θ and f are fixed parameters. In this model, the frequency parameter f is of key interest to our study. We do not perform a fully Bayesian inference by imposing a prior distribution on f because the likelihood function of f is highly irregular, with numerous local maxima, and Monte Carlo computation of the posterior is expensive and intractable for large astronomical surveys.

Previously, [Baluev \(2013\)](#) applied a Gaussian process model to study the impact of red noise in radial velocity planet searches. While his maximum likelihood method is a classical frequentist approach in statistics, our approach can be considered as a hybrid of Bayesian and frequentist approaches. We treat the parameter of interest f , and the parameters for the kernel θ of the Gaussian process as fixed, and impose a prior distribution on other parameters. This is similar to the type-II maximum likelihood estimation of parameters of a Gaussian process or regularization parameters in function estimation; see §5.2 of [Rasmussen & Williams \(2005\)](#). From the Bayesian point-of-view, θ and f are treated as hyper-parameters that in turn are estimated by the empirical Bayes method. Because the Gaussian process plays a critical role in modeling departure of light curves from periodicity, we may also refer to our model more precisely as the non-linear semi-parametric Gaussian process model.

3.2. Estimation of the frequency and the periodogram

Let $\mathbf{y} = (y_1, y_2, \dots, y_n)$ be the observation vector of magnitude of a light curve. By integrating out m, β and \mathbf{h} from the joint distribution given by Eqn. 9, we get the marginal distribution of \mathbf{y} , $p(\mathbf{y}|\theta, f)$, which is a multivariate normal with mean $\mu = m_0 \mathbf{1}$ and covariance matrix

$$\mathbf{K}_y = (\sigma_m^2 + \sigma_b^2 \mathbf{b}_f(t_i)^T \mathbf{b}_f(t_j) + k_\theta(t_i, t_j) + \sigma_i^2 \delta_{ij})_{n \times n},$$

where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ if $i \neq j$. Therefore, the log likelihood of θ and f is

$$\begin{aligned} Q(\theta, f) &= \log(p(\mathbf{y}|\theta, f)) \\ &= -\frac{1}{2}(\mathbf{y} - m_0 \mathbf{1})^T \mathbf{K}_y^{-1} (\mathbf{y} - m_0 \mathbf{1}) \\ &\quad - \frac{1}{2} \log \det \mathbf{K}_y - \frac{n}{2} \log(2\pi). \end{aligned} \quad (10)$$

The maximum likelihood estimator of θ and f is obtained by maximizing $Q(\theta, f)$. Since the likelihood function is differentiable with respect to θ but highly multimodal in the parameter f , standard optimization methods can not be directly used to jointly maximize over θ and f .

We adopt a profile likelihood method as follows. For each frequency f over a dense grid, we compute the maximum likelihood estimator $\hat{\theta}_f = \arg \max_{\theta} Q(\theta, f)$. This can be done using the quasi-Newton method. Then we




Figure 3. The three dimensional surface plot of $Q(\theta, f)$ in Eqn. 10, for the simulated light curve in Fig. 7. Notice $Q(\theta, f)$ is plotted as a function of θ and f is fixed at its true frequency.

perform a grid search to find the maximum profile likelihood estimator of f , i.e.,

$$\hat{f} = \arg \max_f Q(\hat{\theta}_f, f), \quad (11)$$

the estimated period is $\hat{P} = 1/\hat{f}$. The details of the algorithm are given in §3.3. The profile log-likelihood as a function of the frequency f is adopted as the *periodogram* for our model,

$$S_{SP}(f) = Q(\hat{\theta}_f, f). \quad (12)$$

It contains the spectral information of the signal. The frequency of the dominant harmonic component is expected to be the location of the peak of this profile likelihood.

3.3. Computation of the periodogram

Now we present the details of computing the profile likelihood. Because $Q(\theta, f)$ is highly multimodal in the frequency parameter f , we follow the commonly-used strategy of optimization through grid search. On the other hand, since $Q(\theta, f)$ is differentiable in parameter θ , the quasi-Newton method can be employed to optimize over θ for fixed f , and obtain the profile likelihood (Eqn. 12). The gradient of the log likelihood (Eqn. 10) with respect to θ_j ($j = 1, 2$) is

$$\frac{\partial}{\partial \theta_j} Q(\theta, f) = \frac{1}{2} \text{tr} \left((\alpha \alpha^T - \mathbf{K}_y^{-1}) \frac{\partial \mathbf{K}_y}{\partial \theta_j} \right)$$

where $\alpha = \mathbf{K}_y^{-1}(\mathbf{y} - m_0 \mathbf{1})$. In general, the objective function for the Gaussian process model is not convex in its kernel parameters θ and global optimization cannot be guaranteed. Fig. 3 shows the surface plot of $Q(\theta, f)$ as a function of θ for one simulated light curve, with f fixed at the true frequency. The surface exhibits unimodality in this case, although it is not convex.

The computation involved in calculating the profile likelihood through quasi-Newton method can be intensive. Since the objective function (Eqn. 10) is non-convex in θ , generally multiple starting points should be attempted to find the global optimizer when applying the quasi-Newton method. In addition, evaluating the objective function and the gradient function requires inversion of the covariance matrix whose computation cost is of the order $O(n^3)$. During each quasi-Newton iteration, these evaluations could be repeated several times because multiple step size might be attempted. To make the computation more challenging, all of the above need to be repeated at hundreds or even thousands of densely gridded f 's per light curve. Furthermore, the method may need to be applied to hundreds of thousands or millions of light curves from large astronomical surveys.

In order to speed up computation over the dense grid of frequency values, we use the result of applying the quasi-Newton method at one frequency value as a warm start for the subsequent frequency value. Specifically, the optimizer $\hat{\theta}_f$ and its approximate inverse Hessian matrix are provided as quantities to start the quasi-Newton iterations for the next frequency value on the dense grid. When the initial point is near the local minimizer and the inverse Hessian matrix is a good approximation to the true Hessian matrix, the quasi-Newton algorithm will converge in superlinear rate; the step size of $\alpha = 1$ will be accepted by the Wolfe descent condition, avoiding evaluation of the objective function multiple times to determine the appropriate step size during each iteration (see Ch. 6 of Nocedal & Wright 2006, for a more rigorous mathematical discussion). We find that warm start can speed up the computation significantly but sometimes we need to restart with random initial values to ensure convergence to the global optimum. The pseudocode provided in the Appendix describes our algorithm.

3.4. Estimation of the signal and its components

After the parameters f and θ are fixed at their maximum likelihood estimates \hat{f} and $\hat{\theta}_{\hat{f}}$, we can perform the inference of the light curve signal $g(t)$ and its components in the standard Bayesian framework. Interested readers may consult Ch. 2 of Rasmussen & Williams (2005) for a detailed discussion of this topic.

Firstly, we could obtain the posterior distribution of $\gamma = (m, \beta^T)$, the parameters for the long run average magnitude and the exactly periodic term. The prior of γ is $\mathcal{N}(\gamma_0, \Sigma_\gamma)$ with $\gamma_0 = (m_0, 0, 0)^T$ and $\Sigma_\gamma = \text{diag}(\sigma_m^2, \sigma_b^2, \sigma_b^2)$. Its posterior distribution is $\gamma|\mathbf{y} \sim \mathcal{N}(\bar{\gamma}, \bar{\Sigma}_\gamma)$ with

$$\begin{aligned}\bar{\gamma} &= (\mathbf{H}^T \mathbf{K}_c^{-1} \mathbf{H} + \Sigma_\gamma^{-1})^{-1} \\ &\quad (\Sigma_\gamma^{-1} \gamma_0 + \mathbf{H}^T \mathbf{K}_c^{-1} \mathbf{y}), \\ \bar{\Sigma}_\gamma &= (\mathbf{H}^T \mathbf{K}_c^{-1} \mathbf{H} + \Sigma_\gamma^{-1})^{-1},\end{aligned}\tag{13}$$

where

$$\mathbf{h}(t) = (1, \mathbf{b}_{\hat{f}}(t)^T)^T, \mathbf{H} = (\mathbf{h}(t_1), \mathbf{h}(t_2), \dots, \mathbf{h}(t_n))^T,$$

and $\mathbf{K}_c = (k_{\hat{\theta}}(t_i, t_j) + \sigma_i^2 \delta_{ij})_{n \times n}$ with \hat{f} and $\hat{\theta}_{\hat{f}}$ plugged in.

Consider the prediction of light curve magnitude at a specific time point t^* . Define the vector $\mathbf{k}^* = (k_{\hat{\theta}}(t^*, t_1), \dots, k_{\hat{\theta}}(t^*, t_n))^T$. Conditional on (\mathbf{y}, γ) , the distribution of $g(t^*)|\mathbf{y}, \gamma$ is a multivariate normal with mean $\mathbf{h}(t^*)^T \gamma + \mathbf{k}_{\theta}(t^*, \mathbf{t}) \mathbf{K}_c^{-1} (\mathbf{y} - \mathbf{H} \gamma)$ and variance $k_{\hat{\theta}}(t^*, t^*) - (\mathbf{k}^*)^T \mathbf{K}_c^{-1} \mathbf{k}^*$. With the posterior distribution of γ given in Eqn. 13, we are able to remove γ from the above conditional distribution of $g(t^*)$. Finally, we get the posterior distribution of the signal at t^* as $g(t^*)|\mathbf{y} \sim \mathcal{N}(\bar{g}^*, \bar{\sigma}_{g^*}^2)$ with

$$\begin{aligned}\bar{g}^* &= \mathbf{h}(t^*)^T \bar{\gamma} + \mathbf{k}(t^*, \mathbf{t}) \mathbf{K}_c^{-1} (\mathbf{y} - \mathbf{H} \bar{\gamma}), \\ \bar{\sigma}_{g^*}^2 &= k_{\hat{\theta}}(t^*, t^*) - (\mathbf{k}^*)^T \mathbf{K}_c^{-1} \mathbf{k}^* + \mathbf{r}^T \bar{\Sigma}_\gamma \mathbf{r},\end{aligned}\tag{14}$$

where $\mathbf{r} = \mathbf{h}(t^*) - \mathbf{H}^T \mathbf{K}_c^{-1} \mathbf{k}^*$.

4. SIMULATION OF M33 LIGHT CURVES

It is not possible to evaluate the period estimation accuracy of our method directly on the M33 data because the “ground truth” is unknown. Instead, we construct a test data set by smoothing the well-sampled OGLE light curves to infer continuous functions, then resample these functions to match the observational patterns of the M33 data, and at last add noise to the light curves. This data set can serve as a testbed for future studies of comparing different period estimation methods. We will now describe the M33 observations and the construction of the test data set. As the whole simulation procedure is a complicated process, we will discuss its components in detail from §4.1 to §4.4. The whole simulation procedure will be summarized in §4.5.

4.1. Characteristics of the M33 observations

Most of the disk of M33 was observed by the DIRECT (Macri et al. 2001) and M33SSS (Pellerin & Macri 2011) projects in the *BVI* bands, with a combined baseline of 7–9 years and a sampling pattern that depends on the exact location within the disk (see Fig. 4). The large area of coverage and long baseline of these observations make them suitable for Mira searches. We use the *I*-band observations to carry out the simulations, as this is the wavelength range where Miras are brightest (out of the three bands used by these projects). Detailed




Figure 4. Observation patterns for 31 fields in M33, labeled as 0, 1, ..., 9, a, b, ..., u. The horizontal axis shows the Julian date (bottom) and the calendar year (top).

descriptions of the M33 observations can be found in the above referenced papers. We use the data products from a new reduction that will be presented in a companion paper (W. Yuan, et al., in prep.). I -band light curves are available for $\sim 2.5 \times 10^5$ stars, with a median of 44 measurements and a maximum of 170). We model the relation between a magnitude measurement m and its uncertainty σ as:

$$\sigma = a(t'_i, F)^{[m-b(t'_i, F)]} + c(t'_i, F), \quad (15)$$

for each observation field F and each observation night t'_i , where $a(t'_i, F)$, $b(t'_i, F)$ and $c(t'_i, F)$ are field- and night-specific constants. There are 31 different fields in total, $F = 0, 1, \dots, 9, a, b, \dots, u$. The parameters are determined via least-squares fitting using all the measurements for the specific field F and night t'_i . Fig. 5 shows the $m - \sigma$ relation for a typical field.

In order to test the SP periodogram we need sparsely-sampled, moderately noisy Mira light curves with known periods. Thus, we characterized the sampling patterns and noise levels of the M33 observations and simulated Mira light curves of known periods using the OGLE observations of these objects in the LMC.




Figure 5. $m - \sigma$ relation for a given night and field within M33. The solid red line is the best-fit relation using the empirical function $\sigma = a^{(m-b)} + c$, with $a = 2.666$, $b = 23.117$, $c = 0.008$.

4.2. Matching the M33 observation pattern

The first step in simulating a Mira light curve is to randomly select a sampling pattern based on the light curve of an actual star in some field F , $\{t'_i\}_{i=1}^n$ with $n \in [10, 170]$. A random time shift s is added, $t_i = t'_i + s$ for $i = 1, 2, \dots, n$. The random shift s follows a uniform distribution over the interval $[0, P_0]$, where P_0 is the true period of the LMC Mira selected during the artificial light curve generation process. This helps to simulate a large number of unique light curves sampled at random phases using the limited number of template light curves.

4.3. The Mira template light curves

The template Mira light curves are obtained by using our SP model to fit the Mira light curves in the LMC, collected by the OGLE project (Soszyński et al. 2009). A total number of 1663 Miras have been observed in I with very high accuracy, excellent phase coverage and a long baseline (the median and mean number of observations are 466 and 602, respectively, with a baseline of ~ 7.5 years for most fields). Because the LMC light curves are densely sampled with high quality, we can adopt a more complicated model to provide a higher fidelity fit. Following §5.4.3 of Rasmussen & Williams (2005), instead of Eqn. 7, the signal light curve $g(t)$ is decomposed into

$$g(t) = m + l(t) + q(t) + h(t), \quad (16)$$

where m is the long run average magnitude, $l(t)$ is the long-term (low-frequency) trend across different cycles, $q(t)$ is the periodic term, and $h(t)$ is small-scale (high-frequency) variability within each cycle. The latter three terms are modeled by the Gaussian process with different kernels. In particular, we use the squared exponential kernel $k_l(t_1, t_2) = \theta_1^2 \exp(-\frac{1}{2} \frac{(t_1 - t_2)^2}{\theta_2^2})$ for

$l(t)$, another squared exponential kernel $k_h(t_1, t_2) = \theta_6^2 \exp(-\frac{1}{2} \frac{(t_1 - t_2)^2}{\theta_7^2})$ for $h(t)$, and lastly a periodic kernel

$$k_q(t_1, t_2) = \theta_3^2 \exp \left(-\frac{1}{2} \frac{(t_1 - t_2)^2}{\theta_4^2} - \frac{2 \sin^2(2\pi f(t_1 - t_2))}{\theta_5^2} \right)$$

for $q(t)$. Note the periodic kernel allows the light curve amplitude to change across cycles. The maximum likelihood method is applied to fit each LMC light curve, fixing f to the OGLE value and solving for the unknown parameters $(\theta_1, \theta_2, \dots, \theta_7)$. Fig. 6 is an illustration of the model fitting result using Eqn 16 based on the same light curve as in Fig. 2. Notice that the more complex model in Fig. 6 is only suitable for a densely sampled light curve.

Once the sampling pattern is chosen, one of the template light curves will be selected according to the luminosity function described in the next subsection. With the selected template, the magnitude of the simulated light curve signal at t'_i with shift s is $g(t'_i + s)$, which is computed with Eqn. 16 in a similar way as Eqn. 14.

4.4. Matching the luminosity function to the M33 observations

While the OGLE observations of LMC Miras are deep enough to detect these objects over their entire range of luminosities, the M33 observations become progressively more incomplete for fainter and redder objects. We derived an empirical completeness function for the M33 observations as follows. We fit the observed luminosity function $\mathcal{F}_0(I)$ using an exponential for $I \in [18.5, 20]$ mag and extrapolated to fainter magnitudes, obtaining $\mathcal{F}_1(I)$. The empirical completeness function is then $\mathcal{C}(I) = \mathcal{F}_1(I)/\mathcal{F}_0(I)$.

We randomly picked $\{t'_i\}_{i=1}^n$ from the M33 light curves. For each $\{t'_i\}_{i=1}^n$, we selected a (LMC-based) template using $\mathcal{C}(I + 6.2)$ as the probability distribution. The value of +6.2 mag accounts for the approximate difference in distance modulus between the LMC and M33. In this way the resulting luminosity function of the simulated light curves is statistically the same as that of the real M33 observations.

4.5. The simulation procedure

With all the components discussed above, we are able to present the whole simulation procedure here. In order to generate one simulated Mira light curve matching the sampling characteristics of the M33 observations, the first step is to randomly select a sampling pattern $\{t'_i\}_{i=1}^n$, and then add a random shift s , $t_i = t'_i + s$, $i = 1, 2, \dots, n$. The second step is to randomly select a template light curve according to the luminosity function, then compute the light curve signal $g(t'_i + s)$ for the




Figure 6. Light curve of a Mira in the LMC observed by OGLE (black points), decomposed following Eqn 16. Top panel: the fitted light curve; second panel: long-term signal, $m + l(t)$; third panel: periodic term, $m + q(t)$; bottom panel: stochastic variations, $m + h(t)$.

selected sampling pattern $\{t'_i\}_{i=1}^n$. The third step is to use the best-fit relations (Eqn. 15) to add photometric noise via

$$y_i = g(t'_i + s) + 6.2 + \sigma_i \epsilon_i,$$

where +6.2 mag is the approximate relative distance modulus, ϵ_i is drawn from $\mathcal{N}(0, 1)$, and σ_i is computed from

$$\sigma_i = a(t'_i, F)^{[g(t_i) + 6.2 - b(t'_i, F)]} + c(t'_i, F).$$

for the selected observation pattern t'_i and field F . Following this procedure, we generate one simulated light curve $\{t'_i, y_i, \sigma_i\}_{i=1}^n$. The procedure is repeated until 10^5 suitable light curves are generated, excluding any with < 10 data points or sampling on < 7 nights.

5. PERFORMANCE EVALUATION

Having generated the test data set, we evaluate the performance of the SP model and compare it with the generalized Lomb-Scargle (LS) model. We choose prior parameters for the SP model of $m_0 = 15.62 + 6.2$,

$\sigma_m = 10$ and $\sigma_b = 1$. The adopted value of m_0 is the average I magnitude of Miras in LMC and once again $+6.2$ is the approximate relative distance modulus between M33 and the LMC. The values of σ_m and σ_b are larger than those derived from the LMC samples in order to make those priors non-informative. Although fitting the SP model is computationally slower than the LS model, we find that our model gives an overall improvement in various metrics. For both methods, the periodograms are computed on a dense frequency grid from $1/2000$ to $1/100$ with a spacing of the order of 10^{-5} . For the generalized Lomb-Scargle method, we chose a spacing of $(0.05/\text{time span})$ or $\sim 2.5 \times 10^{-5}$, which results in optimal performance for this simulation. For our semi-parametric method, we chose a slightly smaller value of 10^{-5} to facilitate the warm start mechanism in our algorithm (see Appendix) given that small changes in frequency result in tiny changes of the objective function.

5.1. The aliasing effect

We fit the entire simulated data set using the SP model. Fig. 7 gives an example of a simulated light curve and its SP periodogram (Eqn. 12). In this example, the true frequency (labeled by the blue dotted line) is successfully recovered.

Aliasing frequencies at $f \pm 1/365$ d affect most periodograms when dealing with sparsely-observed astronomical data. The red dashed line in Fig. 7 indicates the aliasing frequency at $f + 1/365$ where a strong peak exists. This is not a rare case, and for some light curves the one-year beat aliasing frequencies have higher log likelihoods than the true frequencies. Fig. 8 compares the recovered and true frequencies for all simulated light curves. Two secondary strips parallel to the main one and offset by ± 0.00274 represent $\hat{f} = f \pm 1/365$, respectively. Other aliasing frequencies, such as $2f$, $3f$, $0.5f$, etc., are also noticeable. Lastly, due to the sampling pattern of some light curves, the side lobes of the main peak can be higher than the central value. These manifest as close parallel strips to the aforementioned features.

5.2. Accuracy assessment

The estimated frequency is considered as correct if $\Delta f = |\hat{f} - f_0| < C_f$ for each light curve. The estimation accuracy for the two methods are summarized in Table 1 for several different values of C_f . We choose $C_f = 2.7 \times 10^{-4}$ to stringently bind the one-to-one strip in Fig. 8. Overall, the SP correctly estimates the period for 69.4% of the light curves, while LS model has an accuracy of 63.6%. The improvement of SP over LS is more evident for C-rich Miras, with about 10% higher accuracy, while the improvement for O-rich Miras is




Figure 7. Simulated light curve (top) and periodogram $S_{SP}(f)$ based on our model (bottom). Error bars are derived from the M33 observations. In the top panel, a large gap in temporal coverage was removed to make the plot compact (also note the different time spans). The lines in the bottom panel correspond to the true (solid blue) and the one-year aliasing (dashed red) frequencies.




Figure 8. Model-recovered vs. true frequencies of the test data set. Features other than the one-to-one line are due to aliasing from one-year beat or harmonic frequencies. The three grey diagonal lines correspond to $\hat{f} = f + 1/365$, $\hat{f} = f$ and $\hat{f} = f - 1/365$ from top to bottom, respectively.

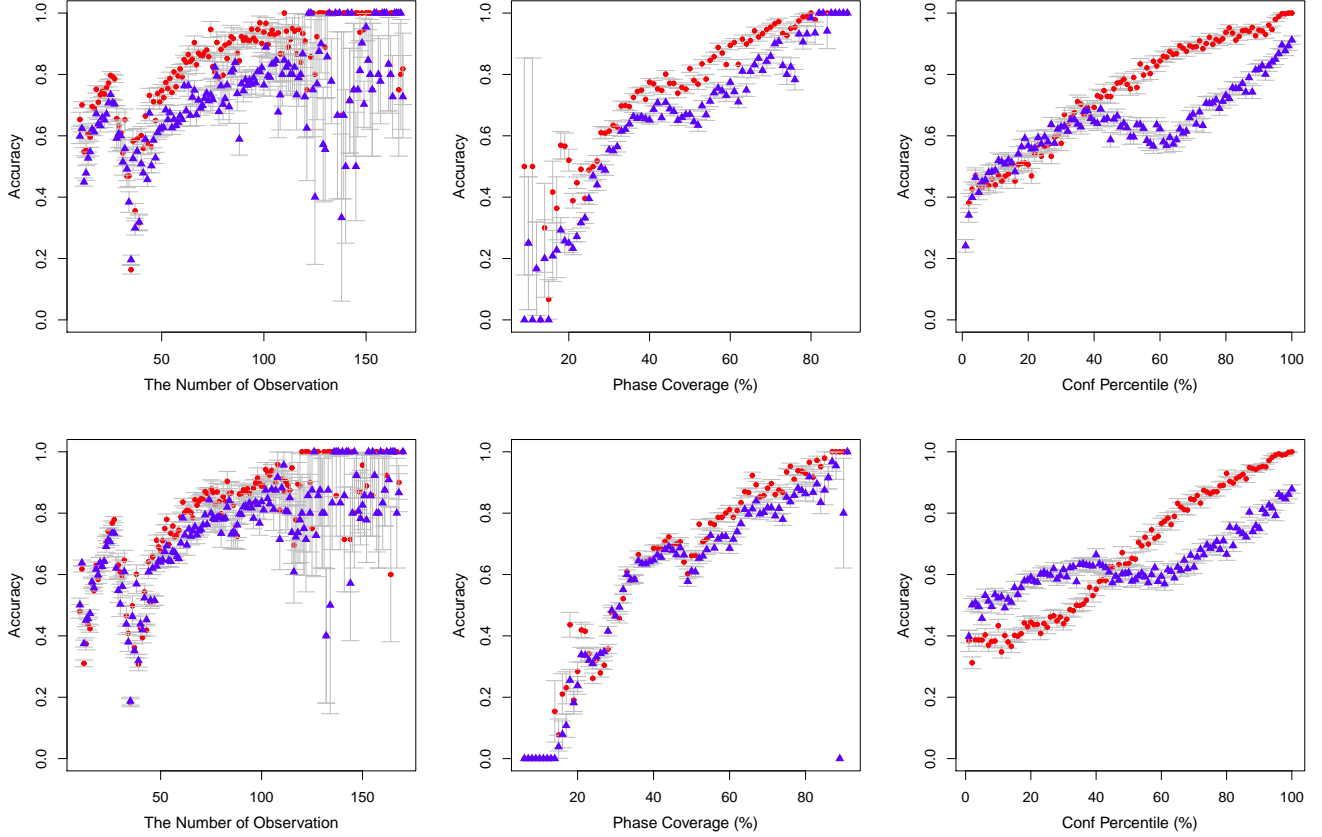


Figure 9. Accuracy comparison using different metrics. Light curves are grouped by number of observations (left), phase coverage (middle), and **conf** values (right). The estimation accuracy in each group is computed and plotted as above. The red circles represent the semi-parametric model (SP), while blue triangles denote the generalized Lomb-Scargle model (LS). The top and bottom panels are for C- and O-rich Miras, respectively.

smaller, with about 3% higher accuracy. The difference in performance arises because C-rich Miras often exhibit larger stochastic deviations that can be better captured by the SP model can better capture, while O-rich Miras have more stable light curves that can be modeled reasonably well with the LS method. We also compute the estimation accuracy of each method by grouping the light curves according to the number of observations, as shown in the left panels of Fig. 9. The top and bottom rows show results for C- and O-rich Miras, respectively. The performance difference is once again more evident in the C-rich category.

Note that accuracy is not a monotonic function of the number of observation, implying this is not a good indicator *per se* of the information content of the light curves for frequency (period) estimation. Thus, we define another metric, called *phase coverage*. Recall that the times of observation for a given light curve are t_1, t_2, \dots, t_n . Given a period of P , these are converted into corresponding phases by $s_i = (t_i \bmod P)/P$, $i = 1, 2, \dots, n$ in the closed interval $[0, 1]$. Now, define

$$J = \left(\bigcup_i (s_i - l, s_i + l) \right) \cap [0, 1],$$

for a specific $l > 0$, the phase coverage can be measured by $\lambda(J)$ where $\lambda(\cdot)$ is the Lebesgue measure (we choose $l = 0.02$). $\lambda(J)$ describes the “length” of the union of the intervals J . For example, $\lambda(J) = 0.1$ for $J = (0.1, 0.2)$, and $\lambda(J) = 0.2$ for $J = (0.1, 0.2) \cup (0.5, 0.6)$.

We divide the light curves into 100 groups such that their $\lambda(J)$ is in one of the interval $(k/100, (k+1)/100]$ for $k = 0, 1, \dots, 99$ and compute the estimation accuracy for each subset. The results for the two models are plotted in the middle column of Fig. 9. Now the estimation accuracy is monotonically increasing as a function of phase coverage. The accuracy improvement of our method is highest when the phase coverage is around 0.5 for C-rich Miras. As the phase coverage approaches the extremes (0 or 1), the performance difference between two methods diminishes. At $\lambda(J) \approx 0$, both methods will fail because this is a hopeless situation. At the other extreme, when $\lambda(J) \approx 1$ and abundant information is available for frequency estimation, both methods have an accuracy close to 1.

The periodogram $S_{SP}(f)$ of our model defined in Eqn. 12 provides more information than just the optimal frequency. Suppose f_1 is the largest local max-




Figure 10. Actual and reconstructed Period-Luminosity relations for C-rich (top) and O-rich (bottom) Miras. The leftmost column shows the actual PLRs using periods and W_I magnitudes from [Soszyński et al. \(2009\)](#). The other columns use the same magnitudes but periods based on the SP or LS algorithms, as indicated in each panel. Recovered PLRs are plotted for various subsets selected according to the `conf` values obtained with the respective method.

imal (global maximum) of $S_{SP}(f)$, and f_2 is the second largest local maximal of $S_{SP}(f)$. Now define $\text{conf} = S_{SP}(f_1) - S_{SP}(f_2) \geq 0$. The value of `conf` serves as a confidence measurement of the global optimal estimate in Eqn. 11. Larger values of `conf` indicate smaller uncertainty in our estimate, and thereby the estimate is more reliable. Now, let c_0 be the smallest value, and let c_1, \dots, c_{100} be the 1th–100th percentiles of all the `conf` values computed for all the light curves. Each light curve can be assigned to a percentile group if its `conf` is in $(c_{k-1}, c_k]$ for some $k \in \{1, 2, \dots, 100\}$. After assigning all light curves by `conf` to their corresponding percentile groups, the estimation accuracy in each group can be computed. The same procedure is applied to the generalized Lomb-Scargle model, with the p -value of the F-statistics given in Eqn. 6 for the top peak being used as its `conf` measurement. The result is plotted in the right column of Fig. 9. The accuracy of our SP method is much higher than the LS model in the top 40 groups. In particular, the accuracy of our method is higher than 90% in the top 20 groups for both C- and O-rich Miras.

Light curves with high values of `conf` are particularly reliable for constructing Period-Luminosity relations (hereafter, PLRs) based on the “Wesenheit” function ([Madore 1982](#)). This function enables a simultaneous correction for the effects of dust attenuation and finite width of the instability strip by defining a new magnitude $W_I = I - 1.55(V - I)$, where V and I are the mean magnitudes in those filters. Figure 10 compares PLRs based on W_I magnitudes and periods determined by OGLE and estimated with each of the two models. The top and bottom rows display the PLRs for C- and O-rich Miras, respectively. The left-most column shows the PLRs based on the actual OGLE periods, while the next two sets of columns show the corresponding relations based on SP or LS periods for the simulated light curves with the top 10 and 40% values of `conf`.

In order to provide a quantitative comparison of the improvement obtained with our SP method, we calculated the dispersion of the actual W_I PLRs and their recovered counterparts as a function of `conf` value as follows, separately for C- and O-rich Miras. First, we

Table 1. Comparison of estimation accuracies (%)

C_f (10^{-4})	Method	Class		
		C-rich	O-rich	Both
1.0	SP	58.1	55.3	56.5
	LS	49.4	51.6	50.6
2.0	SP	69.6	63.8	66.3
	LS	60.1	60.6	60.4
2.7	SP	73.5	66.2	69.4
	LS	63.7	63.5	63.6
N light curves		43,116	56,884	10^5

NOTE—Comparison of estimation accuracy for the semi-parametric (SP) and the Generalized Lomb-Scargle (LS) models. The accuracy is computed for several values of C_f and estimated separately for C- and O-rich Miras; overall values are also given. The number of simulated light curves for each class is also listed.

selected all objects of a given class with $2 < \log P < 3$. If the Soszyński et al. (2009) catalog did not provide a V measurement for a given variable, the missing value was estimated through linear interpolation of the $(I, V-I)$ relation for objects of the same class within $|\Delta \log P| < 0.05$ dex. We fit a quadratic PLR,

$$m = a + b(\log P - 2.3) + c(\log P - 2.3)^2$$

with iterative 3σ clipping (removing $\sim 5\%$ of the data). We then computed the dispersion of the initially-selected OGLE sample about the best-fit relation, including outliers. This yielded “benchmark” dispersions of 0.45 & 0.54 mag for C- & O-rich variables, respectively. Keeping the best-fit relation fixed, we computed the dispersion of recovered PLRs using all artificial light curves within a certain range of *conf* (top 10%, top 20%, ...), using the periods and *conf* values derived by the SP or the LS method. As in the case of the OGLE samples, we only considered objects with $2 < \log P < 3$. The results are plotted in Fig. 11. The SP subsamples exhibit lower (or at worst, equal) dispersions than their LS counterparts for all percentiles and for both subtypes. As discussed previously, the improvement provided by our method is strongest for C-rich Miras and diminishes in significance as one includes light curves with progressively lower confidence values.

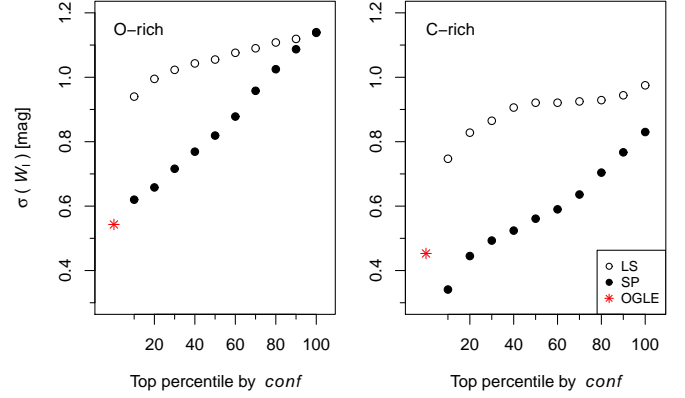


Figure 11. Dispersion of reconstructed Wesenheit P-L relations for different sets of artificial light curves, based on their *conf* value. Top: O-rich PLRs. Bottom: C-rich PLRs. The starred symbols shows the dispersion of the actual OGLE periods and W_I magnitudes. The filled symbols show the dispersion of the recovered PLRs based on SP-derived periods. The open symbols shown the corresponding values for LS-derived periods.

6. SUMMARY

In this paper, we developed a non-linear semi-parametric Gaussian process model for estimating the periods of sparsely-sampled quasi-periodic light curves, motivated by the desire to detect Miras in an existing set of observations of M33. We conducted a large-scale high-fidelity simulation of Mira light curves as observed by the DIRECT/M33SSS surveys to compare our model with the generalized Lomb-Scargle method. Our model shows improved accuracy under various metrics. The simulation data set is provided as a testbed for future comparison with other methods. The semi-parametric model will be used in a companion paper to search for Miras in M33, estimate their periods, and study the resulting period-luminosity relations.

S.H. was partially supported by Texas A&M University-NSFC Joint Research Program. W.Y. and L.M.M. acknowledge financial support from the NSF through AST grant #1211603 and from the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University. J.Z.H. was partially supported by NSF grant DMS-1208952. The authors acknowledge the Texas A&M University Brazos HPC cluster that contributed to the research reported here.

REFERENCES

- Baluev, R. V. 2013, MNRAS, 429, 2052
 Bretthorst, G. L. 1988, Lecture Notes in Statistics, Vol. 48, Bayesian spectrum analysis and parameter estimation (Springer New York)
 Campbell, L. 1925, Harvard College Observatory Reprints, 21, 1
 Deeming, T. J. 1975, Ap&SS, 36, 137
 Glass, I. S., & Evans, T. L. 2003, MNRAS, 343, 67
 Hall, P., Reimann, J., & Rice, J. 2000, Biometrika, 87, 545

- Ita, Y., Tanabé, T., Matsunaga, N., et al. 2004, MNRAS, 353, 705
- Kimeldorf, G., & Wahba, G. 1971, Journal of mathematical analysis and applications, 33, 82
- Lafler, J., & Kinman, T. D. 1965, ApJS, 11, 216
- Liu, D., Lin, X., & Ghosh, D. 2007, Biometrics, 63, 1079
- Lomb, N. R. 1976, Ap&SS, 39, 447
- Ludendorff, H. 1928, Handbuch der Astrophysik, 6, 49
- Macri, L. M., Ngeow, C.-C., Kanbur, S. M., Mahzooni, S., & Smitka, M. T. 2015, AJ, 149, 117
- Macri, L. M., Stanek, K. Z., Sasselov, D. D., Krokenberger, M., & Kaluzny, J. 2001, AJ, 121, 861
- Madore, B. F. 1982, ApJ, 253, 575
- Marsakova, V. I. 1999, Journal of the American Association of Variable Star Observers (JAAVSO), 27, 141
- Nocedal, J., & Wright, S. 2006, Numerical optimization (Springer Science & Business Media)
- O’Sullivan, F., Yandell, B., & Raynor Jr, W. 1986, Journal of the American Statistical Association, 81, 96
- Pellerin, A., & Macri, L. M. 2011, ApJS, 193, 26
- Quinn, B., & Thomson, P. 1991, Biometrika, 78, 65
- Rasmussen, C. E., & Williams, C. K. I. 2005, Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning) (The MIT Press)
- Reimann, J. D. 1994, PhD thesis, University of California, Berkeley
- Scargle, J. D. 1982, ApJ, 263, 835
- Schwarzenberg-Czerny, A. 1996, ApJL, 460, L107
- Soszynski, I., Dziembowski, W. A., Udalski, A., et al. 2007, AcA, 57, 201
- Soszyński, I., Udalski, A., Szymański, M. K., et al. 2009, AcA, 59, 239
- Templeton, M. R., & Karovska, M. 2009, ApJ, 691, 1470
- Templeton, M. R., Mattei, J. A., & Willson, L. A. 2005, AJ, 130, 776
- Udalski, A., Szymanski, M. K., Soszynski, I., & Poleski, R. 2008, AcA, 58, 69
- Wang, Y., Khardon, R., & Protopapas, P. 2012, ApJ, 756, 67
- Wood, P. R., Alcock, C., Allsman, R. A., et al. 1999, in IAU Symposium, Vol. 191, Asymptotic Giant Branch Stars, ed. T. Le Bertre, A. Lebre, & C. Waelkens, 151
- Zechmeister, M., & Kürster, M. 2009, A&A, 496, 577

APPENDIX

Simulated light curves: 10^5 simulated light curves are generated following the procedure of §4. Each light curve is stored in one file with three columns: MJD, I magnitude and uncertainty. The file name, e.g., lc006788.dat is generated sequentially and is only meant for bookkeeping purposes. A mapping between simulated light curve ID and the original OGLE object is given in the file “lc.dat”

Mira variables: This table summarizes the relevant properties of OGLE LMC Miras from Soszyński et al. (2009) that were used to simulate the light curves: OGLE ID, main period and mean I & V magnitudes. It includes some extrapolated values of V for objects with missing data (suitably identified with a “*”). This table can be used to compare true versus derived periods and to generate Period-Luminosity relations.

Software: The related software package, **varStar**, has been released under a GPL3 license (doi.org/10.5281/zenodo.154628). The active software development repository can be found at github.com/shiuyanhe/varStar.

PSEUDO-CODE

procedure QUASINEWTON() *Quasi-Newton’s Method with Grid Search*

Input: Maximal and minimal trial frequencies $f_M > f_m > 0$; frequency step Δf ; n observations $\{t_i, y_i, \sigma_i\}$.

Output: Periodogram $S(f)$ evaluated at the trial frequencies.

- 1: Initialize $\theta^{(0)}$ and $\mathbf{H}^{(0)}$, and $f \leftarrow f_m$;
- 2: **for** $f \in \{f_m, f_m + \Delta f, f_m + 2\Delta f, \dots, f_M\}$ **do**
- 3: $p \leftarrow 0$;
- 4: **repeat**
- 5: $\mathbf{t}_p \leftarrow -\mathbf{H}^{(p)} \frac{\partial}{\partial \theta} Q(\theta^{(p)}, f)$;
- 6: $\theta^{(p+1)} \leftarrow \theta^{(p)} + \alpha_p \mathbf{t}_p$ and the step size α_p satisfying the Wolfe condition;
- 7: $\mathbf{d}_p \leftarrow \theta^{(p+1)} - \theta^{(p)}$, $\mathbf{e}_p \leftarrow \frac{\partial}{\partial \theta} Q(\theta^{(p+1)}, f) - \frac{\partial}{\partial \theta} Q(\theta^{(p)}, f)$;
- 8: $\rho_p = 1/\mathbf{d}_p^T \mathbf{e}_p$;
- 9: $\mathbf{H}^{(p+1)} \leftarrow (\mathbf{I} - \rho_p \mathbf{d}_p \mathbf{e}_p^T) \mathbf{H}^{(p)} (\mathbf{I} - \rho_p \mathbf{e}_p \mathbf{d}_p^T) + \rho_p \mathbf{d}_p \mathbf{d}_p^T$; ▷ BFGS update
- 10: $p \leftarrow p + 1$;
- 11: **until** $\|\mathbf{d}_{p-1}\| < \epsilon$
- 12: $\hat{\theta}_f \leftarrow \theta^{(p)}$, and $S(f) \leftarrow Q(\hat{\theta}_f, f)$;
- 13: $\mathbf{H}^{(0)} \leftarrow \mathbf{H}^{(p)}$ and $\theta^{(0)} \leftarrow \theta^{(p)}$; ▷ Save warm start value for next trial frequency
- 14: **end for**